

# Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-01 19:43:30

PAGE 1

REFERENCE NO: 202

This contribution was submitted to the National Science Foundation as part of the NSF CI 2030 planning activity through an NSF Request for Information, [https://www.nsf.gov/publications/pub\\_summ.jsp?ods\\_key=nsf17031](https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf17031). Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

## Author Names & Affiliations

- Nikolai Pogorelov - Department of Space Science, University of Alabama in Huntsville

## Contact Email Address (for NSF use only)

(Hidden)

## Research Domain, discipline, and sub-discipline

plasma physics, space physics, solar and heliospheric physics; applied mathematics and programming.

## Title of Submission

On the Need of the Cyberinfrastructure Development Led by the NSF

## Abstract (maximum ~200 words).

The National Science Foundation was created to "promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense." Most of the fundamental science projects are supported by the NSF. The success of its mission strongly depends on the availability of proper experimental facilities, computing hardware and software, and the workforce capable of using all of them efficiently. In the era where science and engineering communities are rapidly becoming digital and computational, and data science is ever more critical to progress, the country needs to follow a strong path for supporting computation- and data-intensive research across a broad spectrum of research disciplines. Everybody involved high-performance computing these days knows that facilities are saturated and that researchers are in need of national data services. In this situation, the scientific community regards NSF as the only agency that can propose a clear, long-term plan for the cyberinfrastructure development that would serve the needs of the broad spectrum of scientists and engineers in their pursue of discoveries in the fundamental science and technology. Here, the challenges are discussed in plasma physics research, in general, and space physics, in particular. This discussion is followed by a brief description of the workforce development in the time when the computational and computer sciences are strongly overlapped.

**Question 1** Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The subject of space physics covers a wide range of topics from the solar interior, solar wind (SW) acceleration near the Sun, Sun–Earth connection, planetary magnetospheres, to the SW interaction with the magnetized local interstellar medium (LISM). Physical problems involved include turbulence, kinetic processes in the collisionless SW plasma, ion acceleration, magnetic reconnection, various space weather topics, non-thermal ion behavior, charge-exchange effects, ionosphere-magnetosphere coupling, the physics of coronal mass ejections, galactic cosmic ray transport throughout the heliosphere, the origin and behavior of anomalous cosmic rays, and many others.

# Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-01 19:43:30

PAGE 2

REFERENCE NO: 202

The range of problems stipulates the necessity to address different space and time scales and use different sets of equations, e.g., MHD, Boltzmann, Fokker–Planck, Parker, focused-transport, and other equations, which very often should be solved together, in a coupled fashion. Numerical methods used to solve those equations are also different and can be based on both deterministic and stochastic approaches. Coupling of different scales frequently requires overlapping grids, which is impossible without adaptive mesh refinement (AMR) on the largest scale. An increase in the AMR level creates challenges for data exchange between computing nodes. Coupling of different sets of equations requires more sophisticated algorithms that take into account computer architectures, e.g., the need in hybrid parallelization strategies.

Although kinetic and hybrid simulations provide invaluable insight into the microscopic nature of such physical phenomena as, e.g., magnetic reconnection, scaling of the obtained results to macroscopic scales may be very difficult, which is especially well seen in the attempts to

explain Voyager 1 observations at the SW–LISM boundary. It is interesting that coupling between microscopic and macroscopic scales occurs through the turbulence intrinsic to collisionless space plasma. The mere possibility to apply the equations of continuum mechanics (gas dynamics or MHD equations) is due to a considerable decrease in the ion mean free path due to their scattering on the magnetic field fluctuations. The presence of turbulence undermines the frozen-in conditions for magnetic field lines and therefore affects magnetic reconnection at a macroscopic level. To couple the scales in global simulations one needs to identify the spectrum of turbulent fluctuations in the regions of possible reconnection and instability. This can be done by performing hybrid (ions are treated as particles while electrons behave as the neutralizing

fluid) simulations on a grid overlapping with the global grid. This is feasible with AMR. On the other hand, hybrid simulations may overlap in scale with full particle simulations that take into account electron kinetic effects. A hierarchy of overlapping meshes makes it possible to couple macro- and micro-scales in many problems. Available computing power does not allow us to cover all scales self-consistently, but the setup for each lower-scale problem may be defined by a one-level-higher-scale problem solution. Additionally, the knowledge of the properties of MHD turbulence determined locally in problem-specific regions makes it possible to describe ion acceleration and GCR transport using realistic transport coefficients, and relevant plasma and magnetic field background. Below is a brief list of our research needs.

1. Global AMR simulations based on MHD equations coupled with the Boltzmann and Fokker–Planck equations are necessary to model the SW–LISM interaction and a variety of related fundamental physical phenomena (ion acceleration, MHD and kinetic instabilities, magnetic reconnection, etc., intrinsic to most problems involving collisionless space plasma and essential for space weather prediction as one of the NSF

priority research subjects) measured in situ by Voyagers and remotely by IBEX. Such simulations belong to the capability computing area. They require efficient system-level support for parallel code-to-code communication, so that two different parallel codes could run simultaneously and be able to communicate with each other. So far we have been dealing with this problem taking advantage of the fact that the kinetic part

requires considerably greater time than its MHD counterpart. In this case the parallelization deficiencies are not visible and it suffices only to arrange an optimum distribution of the MHD data among the cores solving the Boltzmann equation. This was done using a mixed MPI/OpenMP parallelization strategy and special data buffering.

2. Substantial computing resources are necessary to perform global, 3D, hybrid simulations of the SW plasma, especially in the Earth magnetosphere. While some results on such modeling have been reported recently, most of hybrid simulations are either local or 2D. These problems are of major importance for understanding the SW interaction with planets, including Earth. There are attempts to embed hybrid simulations into AMR MHD codes, which saves computer resources (see, however, the previous paragraph).

3. Full-particle (particle-in-cell, PIC) simulations are necessarily limited to very small regions of the order of a few tens of ion gyroradii. These simulations are among the most computationally demanding calculations that are run on the NSF and DOE supercomputers. PIC codes are easily parallelized, but are data and computationally expensive. They are important for phenomena dependent on the electron dynamics. An implicit approach has been proposed recently for PIC simulations, which however essentially smears electron scales. It is necessary to acknowledge in this connection that PIC and MHD scale do not overlap, which undermines attempts to combine these two approaches in one code highly challenging. No wonder, such efforts (currently supported by the NSF-DOE program) are focused on combining an MHD code with an implicit PIC code.

4. The physics of fusion and space physics share a fundamental plasma property: the presence of both thermal and non-thermal ions. Charge exchange between the interstellar atoms and solar wind ions is a kinetic process that should be analyzed by solving the Boltzmann equation. So is the charge exchange that is a fundamental component in the neutral beam heating of plasma. Both communities solve the Fokker-Planck (FP) equations of different level of complexity to analyze the behavior of non-thermal ions. In space physics, this

# Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-01 19:43:30

PAGE 3

REFERENCE NO: 202

phenomenon is complicated by the presence of shocks. A challenge is to understand the microphysics of pickup ions crossing such shocks and build accurate, coupled MHD-FP models based on it. In addition, both communities require understanding of electron acceleration by instabilities in the pickup ion distribution and by shocks propagating through plasma, and of Langmuir waves launched by the accelerated electrons.

5. Magnetic reconnection is a fundamental problem both for plasma physics and in space physics. The major difficulty is that sometimes we need to model macroscopic effects of magnetic reconnection, while the physics of this phenomenon is based on microscopic effects. Fortunately, there is a strong indication that pre-existing and self-generated turbulence, e.g. near the Sun, may be more important for magnetic reconnection than the microphysics. This opens new possibilities for modeling phenomena based on magnetic reconnection.

6. Solar Probe Plus (SPP) is a new NASA mission with the launch scheduled in 2018. Among other high-resolution data, this mission will perform time-dependent measurements of the 3D, ion, electron and  $\alpha$ -particle distribution functions. On the one hand, this will result in an enormous amount of data, considering the 7-dimensional character of the distribution. On the other hand, this will require the development of new, efficient numerical methods to interpret the observations. This means the necessity of developing 6-dimensional parallelization strategies and AMR both in the physical and velocity space. The SPP mission will give rise to a data-intensive simulations in space physics.

**Question 2** Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

The research challenges outlined in the previous section require HPC. The following cyberinfrastructure issues should therefore be addressed.

1. To make computational efforts more efficient, the NSF should prioritize funding research efforts that focus on new intelligent numerical algorithms, i.e., asynchronous time integration, AMR, load-balancing, etc.

2. Efficient, parallel visualization software for analyzing multi-dimensional numerical results should be supported and made accessible on every supercomputer.

3. An area that is poorly covered by existing systems is data storage and sharing. Most large-scale computational centers do have adequate storage facilities, but none of them provides reasonable resources for collaborative data analysis and data sharing, while many large-scale simulations generate data of interest to many research groups.

5. High-performance resources for the space physics community are available on NASA, NSF, and DOE supercomputers. However, computing power available for scientific simulation on NASA supercomputers is rather limited and, of course, is strongly mission-oriented. On the other hand, the DOE INCITE program imposes requirement on the code scalability so strong that getting allocations on Titan or Mira is very difficult for most of users. In my opinion, NSF allocations are essential for the majority of simulation-related NSF projects. The demand in computing resources will continue to exceed available capabilities. The progress on the hardware side greatly outpaces the progress in the algorithms and software development, which is underfunded across the agencies. It is desirable for the NSF to support the development of a few high-quality community codes and for the space physics community to take advantage of the already developed software.

**Question 3** Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Clearly, the development and maintaining of the cybercomputing infrastructure by NSF should be accompanied by the development of workforce capable of using new software and hardware technologies. In most of cases, this is not the case. The codes are written by a

# Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-01 19:43:30

PAGE 4

REFERENCE NO: 202

---

relatively small group of physicists who accept the burden to learn from computer scientists. The code I have developed with my colleagues (Multi-Scale Fluid-Kinetic Simulations Suite, MS-FLUKSS) was written on top of the Chombo adaptive mesh-refinement framework. This helped us to save efforts, but still required 4 years of intensive work supported by the NSF ITR program. My graduate students at UAH usually do not have enough experience to work with such codes, so I encourage them to join a number of courses offered by the UAH Department of Computer Science. In my opinion, NSF should offer fellowships for students interested both in physics and programming. For example, one of my graduate students was a Blue Waters fellow, but this fellowship lasts only one year.

Students that have experience in writing sophisticated contemporary codes have very high chances to be hired by rich private companies, where they will work on new technologies. However, considerable progress should be expected even if a small fraction of them remains in science.

The scientific and engineering community in the country is not passively watching the apparent decrease in the HPC availability and accessibility. In January of 2015, I participated in the Brainstorming HPCD Workshop in Dallas, TX, which was part of a community-driven effort to gather comprehensive information about the needs of the NSF-sponsored science and engineering research community for advanced capacity through capability computing, including data-intensive and compute-intensive resources. The workshop focused on the research community needs at the high end of the spectrum and attempted to articulate the value, contributions and impacts of at-scale computing for the advancement of science, engineering, and other research. The final report of this workshop provided input to the National Research Council Committee on Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science in 2017-2020. The final report is available on <http://brainstormhpcd.org/final-report/>. In addition, a group of enthusiastic researchers, including myself, got together in Washington, DC in August, 2016, where a document describing the urgent need of NSF-driven program of HPC development was prepared and submitted to the NSF's Directorate for Computer and Information Science and Engineering.

In summary, the scientific community all over the country has strong hopes that NSF will lead the HPC efforts in the interest of researchers ranging from physicists to drug developers in the NIH institutions. On the other hand, NSF should be assured that our community itself will provide any help necessary in the promotion of modern cyberinfrastructure and preparation of workforce that will fit new related challenges.

## Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>)."
-